arXiv:1602.04941v1 [math.ST] 16 Feb 2016

On Perfect Classification for Gaussian Processes

Juan A. Cuesta-Albertos[1] and Subhajit Dutta

Departamento de Matemáticas, Estadística y Computación, Universidad de Cantabria, Spain.
Department of Mathematics and Statistics, IIT Kanpur, India.

February 17, 2016

Abstract

In this paper, we study the problem of discriminating J (≥ 2) Gaussian processes by analyzing the behavior of the underlying probability measures in an infinite-dimensional space. Motivated by singularity of a certain class of Gaussian measures, we first propose a data based transformation for the training data. For a J class classification problem, this transformation induces complete separation among the associated Gaussian processes. The misclassification probability of a componentwise classifier when applied on this transformed data asymptotically converges to zero. In finite samples, the empirical classifier is constructed and related theoretical properties are studied. Good performance of the proposed methodology is demonstrated using simulated as well as benchmark data sets when compared with some parametric and nonparametric classifiers for such functional data.

Some key words: Bayes' risk, Consistency in probability, Cross-validation, Difference in covariance operators, Hajek and Feldman property, Mahalanobis' distances.

Short title: Discrimination of Gaussian processes

[1]This author has been partially supported by the Spanish Ministerio de Ciencia y Tecnología, grants MTM2011-28657-C02-02 and MTM2014-56235-C2-2-P.

1 Introduction

Suppose that we are given two Gaussian processes (GPs) $\mathbf{X}$ and $\mathbf{Y}$ defined on a common probability space $(\Omega, \sigma, \mathbb{P})$, and the probability distributions that they generate are $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$, respectively. The Hajek and Feldman property (established independently in [8] and [5]) states that $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$ are either equivalent, or else mutually singular. In other words, for every measurable set A, $\mathbb{P}_{\mathbf{X}}(A) = 0$ if and only if (iff) $\mathbb{P}_{\mathbf{Y}}(A) = 0$, or else there exist two measurable sets $S_{\mathbf{X}}$ and $S_{\mathbf{Y}}$ such that

$$\mathbb{P}_{\mathbf{X}}(S_{\mathbf{X}}) = 1, \mathbb{P}_{\mathbf{Y}}(S_{\mathbf{X}}) = 0 \text{ and } \mathbb{P}_{\mathbf{X}}(S_{\mathbf{Y}}) = 0, \mathbb{P}_{\mathbf{Y}}(S_{\mathbf{Y}}) = 1.$$

Some work has been done to find conditions under which equivalence, or singularity between two GPs occur (see, e.g., [12], [13], [14], [15], [16]).

More recently, the nice paper [3] analyzes the classification problem in which an observation $\mathbf{z}$ is to be classified as produced either from $\mathbb{P}_{\mathbf{X}}$ or $\mathbb{P}_{\mathbf{Y}}$. This paper investigates conditions under which a perfect classification procedure is possible. As stated above, this is only possible if $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$ are mutually singular. To the best of our knowledge, this paper is the first time in which the Hajek and Feldman property has been applied in practice to solve a statistical problem. The relationship between the paper [3], and this property has been analyzed in [2], where the authors also present expression of the optimal Bayes' rule in some classification problems. Bayes' classification for GPs was also studied in [1], and the authors used the equivalence of certain GPs to derive a classification procedure based on the likelihood ratio. However, we cannot write the likelihood ratio if the GPs are singular.

The work in [3] is devoted to the case where $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$ differ in locations. We propose a procedure which allows one to identify mutually singular situations based not only on locations, but also when there is difference in the covariance operators. The procedure is based on analyzing the limit of some Mahalanobis' distances defined on suitable sequences of increasing finite-dimensional spaces. One of the main differences between the results stated in [3], and those we present here is that [3] includes a result in which they identify the optimal projection to be used to separate $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$, assuming the covariance operators to coincide. Here, we are considering a kind of functional Behrens-Fisher problem [17], where the optimal solutions does not exist even in the real case. We want to mention that a classification procedure for functional data based on Mahalanobis' distances is proposed recently in [6], but 'perfect classification' is not achieved because their procedure stops at a finite-dimensional subspace.

The structure of this paper continues as follows. In Section 2, we first propose a transformation and analyze the related classification problem when the distributions are known. The notation to be used in the theoretical results are introduced in Subsection 2.1. This section ends by explaining the application of those results to the classification problem. In Section 3, we examine some examples for application of our results and compare them with the results stated in [3]. Extension of our method to the case where the parent distributions are unknown, and need to be estimated from the training data is carried out in Section 4. Sections 5 and 6 are devoted to the analysis of the behavior of

the procedure in practice; the first of those sections contains simulations, while the latter presents results of our method applied to some benchmark data sets. The paper ends with a technical Appendix which contains proofs of the theorems, and several technical results.

2 Classification with Known Distributions

Assume that the processes $\mathbf{X}$ and $\mathbf{Y}$ are defined on a bounded real interval, which without loss of generality, we identify with the unit interval $[0,1]$. We will also assume that they belong to the Hilbert space of square integrable functions:

$$\mathbb{H}: \text{ the set of the real functions } f(t) \text{ with } t \in [0,1] \text{ such that } \int_0^1 f^2(t)dt < \infty.$$

Here, the inner product is defined as $\langle f, g\rangle = \int_0^1 f(t)g(t)dt$, and the norm is defined as $\|f\| = \left(\int_0^1 f^2(t)dt\right)^{1/2}$. To avoid repetitions, we use $\mathbf{Z}$ to indistinctly denote both processes $\mathbf{X}$ and $\mathbf{Y}$. The distribution of the random process $\mathbf{Z}$ will be denoted as $\mathbb{P}_{\mathbf{Z}}$, its mean function by $\mu^{\mathbf{Z}}$ and its covariance by $\Sigma^{\mathbf{Z}}$. We will also assume that all involved random quantities are defined on the common and rich enough probability space $(\Omega, \sigma, \mathbb{P})$. If we have an observation $\mathbf{z} \in \mathbb{H}$ taken either from $\mathbb{P}_{\mathbf{X}}$ or $\mathbb{P}_{\mathbf{Y}}$, our aim is to discriminate between those two possibilities. In this section, the distributions $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$ are assumed to be known.

Consider an increasing sequence of subspaces $\{V_d\}_{d=1}^{\infty}$, with $V_d \subset \mathbb{H}$, where the dimension of V_d is d for $d \in \mathbb{N}$. This restriction is not necessary for the development which follows as long as the dimension of V_d goes to infinity with d, but it simplifies the notation. Given the subspace V_d, let $\mu_d^{\mathbf{Z}}$ denote the d-dimensional mean of the projection of $\mathbf{Z}$ on V_d, and $\Sigma_d^{\mathbf{Z}}$ represent the $d \times d$-dimensional covariance matrix of the projection of $\mathbf{Z}$ on this subspace. For every $d \in \mathbb{N}$, we denote the eigenvalues of the matrix $\Sigma_d^{\mathbf{Z}}$ in decreasing order by $\lambda_{d,i}^{\mathbf{Z}}$ for $i = 1, \ldots, d$. In general, the fact that $V_d \subset V_{d+1}$ does not guarantee the existence of any relationship between the sets $\{\lambda_{d,1}^{\mathbf{Z}}, \ldots, \lambda_{d,d}^{\mathbf{Z}}\}$ and $\{\lambda_{d+1,1}^{\mathbf{Z}}, \ldots, \lambda_{d+1,d+1}^{\mathbf{Z}}\}$. In Section 4, we will assume that V_d is generated by the first d eigenvectors of the process $\mathbf{Z}$. In that case, it trivially happens that

$$\lambda_{d,i}^{\mathbf{Z}} = \lambda_{i,i}^{\mathbf{Z}} \text{ for } d \geq i, \text{ with } i \in \mathbb{N}.$$

To avoid technical problems, we will assume that all the eigenvalues of the covariance operators of $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$ are strictly positive and different. Thus, $\lambda_{d,i}^{\mathbf{Z}} > 0$ for every i with $1 \leq i \leq d$.

If $\mathbf{u} \in \mathbb{H}$, we denote $\mathbf{u}_d$ to be its projection on V_d and write $\mathbf{u}_d = (u_{d,1}^{\mathbf{Z}}, \ldots, u_{d,d}^{\mathbf{Z}})^T$, when we take in V_d the basis composed by the eigenvectors of $\Sigma_d^{\mathbf{Z}}$. Given $d = 1, \ldots$ and $\mathbf{u} \in \mathbb{H}$, the square of the Mahalanobis norm of $\mathbf{u}_d$ associated with the covariance matrix

$\Sigma_d^{\mathbf{Z}}$ is as follows:

$$\left\|(\Sigma_d^{\mathbf{Z}})^{-1/2}\mathbf{u}_d\right\|^2 = \sum_{i=1}^{d} \frac{\left(u_{d,i}^{\mathbf{Z}}\right)^2}{\lambda_{d,i}^{\mathbf{Z}}}.$$

Here, $\|\cdot\|$ denotes the usual Euclidean norm. Given $\mathbf{u}, \mathbf{v} \in \mathbb{H}$, a keystone of this paper is the following map:

$$D_d^{\mathbf{Z}}(\mathbf{u}, \mathbf{v}) = \frac{1}{d}\|(\Sigma_d^{\mathbf{Z}})^{-1/2}(\mathbf{u} - \mathbf{v})_d\|^2. \tag{1}$$

2.1 The classification procedure

The classification procedure that we propose is based on the behavior of the statistics $D_d^{\mathbf{X}}(\mathbf{z}, \mu^{\mathbf{X}})$ and $D_d^{\mathbf{Y}}(\mathbf{z}, \mu^{\mathbf{Y}})$ depending on the distribution of $\mathbf{Z}$ (either $\mathbb{P}_{\mathbf{X}}$ or $\mathbb{P}_{\mathbf{Y}}$). This analysis is carried out in Theorem 2.1 below. The theorem is based on a more general Proposition 7.1 which is presented in the Appendix.

Theorem 2.1 *Assume that* $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$. *Then,*

$$D_d^{\mathbf{X}}(\mathbf{z}, \mu^{\mathbf{X}}) \xrightarrow{P} 1 \text{ as } d \to \infty. \tag{2}$$

Let $\alpha_1^d, \ldots, \alpha_d^d$ *denote the eigenvalues of the matrix* $S_d = (\Sigma_d^{\mathbf{Y}})^{-1/2}\Sigma_d^{\mathbf{X}}\,(\Sigma_d^{\mathbf{Y}})^{-1/2}$. *Moreover, if there exist constants* $L_\mu^{\mathbf{X},\mathbf{Y}}$ *and* $L_S^{\mathbf{X},\mathbf{Y}}$ *(finite, or not) such that*

$$L_\mu^{\mathbf{X},\mathbf{Y}} = \lim_d \frac{1}{d}\left\|(\Sigma_d^{\mathbf{Y}})^{-1/2}(\mu^{\mathbf{Y}} - \mu^{\mathbf{X}})_d\right\|^2, \tag{3}$$

$$L_S^{\mathbf{X},\mathbf{Y}} = \lim_d \frac{1}{d}\, trace(S_d), \text{ and} \tag{4}$$

$$0 = \lim_d \frac{\sup(\alpha_1^d, \ldots, \alpha_d^d)}{d}, \tag{5}$$

then

$$D_d^{\mathbf{Y}}(\mathbf{z}, \mu^{\mathbf{Y}}) \xrightarrow{P} (L_\mu^{\mathbf{X},\mathbf{Y}} + L_S^{\mathbf{X},\mathbf{Y}}) \text{ as } d \to \infty. \tag{6}$$

Proof : Statement (2) follows by taking $\mu = \mu^{\mathbf{X}}$, and $A_d = \Sigma_d^{\mathbf{X}}$ in Proposition 7.1. We obtain (6) by taking $\mu = \mu^{\mathbf{Y}}$, and $A_d = \Sigma_d^{\mathbf{Y}}$ in the same proposition. •

Remark 2.1.1 A condition in Theorem 2.1 is required to ensure that no single component is extremely influential. For instance, it may happen that we take a sequence such that $\alpha_d^1 = d$ and $\alpha_d^i = o(d^{-1})$ for every $i \geq 2$. Under this condition, no limit is possible in Theorem 2.1 (see the proof of Proposition 7.1). This possibility is excluded by assumption (5).

Remark 2.1.2 Assume that the subspace V_d is generated by the first d eigenfunctions of $\Sigma^{\mathbf{X}}$. Lemma 7.2 (stated in the Appendix) shows that assumption (5) in Theorem 2.1 follows from assumption (4) when $L_S^{\mathbf{X},\mathbf{Y}}$ is finite. So, it is sufficient to check assumption (4) in the examples with GPs that we study later in Section 3.

Remark 2.1.3 It is obvious that Theorem 2.1 holds if we replace $\mathbb{P}_{\mathbf{X}}$ by $\mathbb{P}_{\mathbf{Y}}$. Then, under the corresponding assumptions for $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{Y}}$, we get

$$D_d^{\mathbf{X}}(\mathbf{z}, \mu^{\mathbf{X}}) \xrightarrow{P} (L_\mu^{\mathbf{Y},\mathbf{X}} + L_S^{\mathbf{Y},\mathbf{X}}) \text{ and } D_d^{\mathbf{Y}}(\mathbf{z}, \mu^{\mathbf{Y}}) \xrightarrow{P} 1 \text{ as } d \to \infty.$$

In order to apply Theorem 2.1 to classification problems, we define a sequence of two-dimensional statistics as follows:

$$\mathbf{T}_d(\mathbf{z}) = (D_d^{\mathbf{X}}(\mathbf{z}, \mu^{\mathbf{X}}), D_d^{\mathbf{Y}}(\mathbf{z}, \mu^{\mathbf{Y}}))^T \text{ for } d \in \mathbb{N}.$$

Based on $\mathbf{T}_d$, let us consider the following classifier

$$\delta_d(\mathbf{z}) = \begin{cases} 1, & \text{if } |D_d^{\mathbf{X}}(\mathbf{z}, \mu^{\mathbf{X}}) - 1| \leq |D_d^{\mathbf{Y}}(\mathbf{z}, \mu^{\mathbf{Y}}) - 1|, \\ 2, & \text{else.} \end{cases}$$

For a fixed value of $d \geq 1$, we define p_d to be the total misclassification probability of the classifier δ_d, i.e., $p_d = \pi_{\mathbf{X}} \mathbb{P}_{\mathbf{X}}[\delta_d = 2] + \pi_{\mathbf{Y}} \mathbb{P}_{\mathbf{Y}}[\delta_d = 1]$, with $\pi_{\mathbf{X}}$ (respectively, $\pi_{\mathbf{Y}}$) being the proportion of observations coming from $\mathbb{P}_{\mathbf{X}}$ (respectively, $\mathbb{P}_{\mathbf{Y}}$). We assume that $\pi_{\mathbf{X}} + \pi_{\mathbf{Y}} = 1$, and $\pi_{\mathbf{X}}, \pi_{\mathbf{Y}} > 0$.

Theorem 2.2 *If $L_\mu^{\mathbf{X},\mathbf{Y}} + L_S^{\mathbf{X},\mathbf{Y}} \neq 1$ and $L_\mu^{\mathbf{Y},\mathbf{X}} + L_S^{\mathbf{Y},\mathbf{X}} \neq 1$, then the misclassification probability $p_d \to 0$ as $d \to \infty$.*

Proof: Consider a fixed $\mathbf{z}$. Under the assumptions mentioned above in Theorem 2.1 and Remark 2.1.3, it happens that for $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$,

$$|D_d^{\mathbf{X}}(\mathbf{z}, \mu^{\mathbf{X}}) - 1| \xrightarrow{P} 0 \text{ as } d \to \infty,$$

while

$$|D_d^{\mathbf{Y}}(\mathbf{z}, \mu^{\mathbf{Y}}) - 1| \xrightarrow{P} |L_\mu^{\mathbf{X},\mathbf{Y}} + L_S^{\mathbf{X},\mathbf{Y}} - 1| \text{ as } d \to \infty.$$

The latter expression equals 0 iff $L_\mu^{\mathbf{X},\mathbf{Y}} + L_S^{\mathbf{X},\mathbf{Y}} = 1$. Thus, the assumption in the statement gives $I_{[\delta_d(\mathbf{Z})=2]} \xrightarrow{P} 0$ as $d \to \infty$. Applying the Dominated Convergence Theorem, for any positive value of $\pi_{\mathbf{X}}$, we now get $\mathbb{P}_{\mathbf{X}}[\delta_d = 2] \to 0$ as $d \to \infty$.

Similarly, when $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{y}}$, the first limit is zero iff $L_\mu^{\mathbf{Y},\mathbf{X}} + L_S^{\mathbf{Y},\mathbf{X}} = 1$. A similar line of arguments continue to hold, and we have $\mathbb{P}_{\mathbf{Y}}[\delta_d = 1] \to 0$ as $d \to \infty$. This completes the proof. •

A consequence of Theorem 2.2 is that we can now identify the distribution which generated $\mathbf{z}$ without a possibility of mistake as $d \to \infty$, i.e., we obtain '*perfect classification*'.

3 Some Example of GPs

We now analyze the limiting behavior of the statistic $\mathbf{T}_d(\mathbf{z})$ in the particular case when $\mu^{\mathbf{X}} = 0$ and $\Sigma^{\mathbf{Y}} = a\Sigma^{\mathbf{X}}$ (here a is a positive constant). Consider V_d to be the subspace generated by the first d eigenfunctions of the covariance operator $\Sigma^{\mathbf{X}}$. Assume that there exists a positive constant ν such that $\nu = \lim_{d\to\infty} \frac{1}{d}\|(\Sigma^d_{\mathbf{X}})^{-1/2}\mu^{\mathbf{Y}}_d\|^2$. This now implies that we have definite expressions for the limiting constants mentioned in equations (3) and (4)

$$L^{\mathbf{X},\mathbf{Y}}_{\mu} = \nu/a, L^{\mathbf{X},\mathbf{Y}}_{S} = 1/a \text{ and } L^{\mathbf{Y},\mathbf{X}}_{\mu} = \nu, L^{\mathbf{Y},\mathbf{X}}_{S} = a.$$

The following simplified limits hold for $\mathbf{T}_d(\mathbf{z})$:

$$\mathbf{T}_d(\mathbf{z}) \xrightarrow{P} \begin{cases} (1,(\nu+1)/a)^T, & \text{if } \mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}, \\ (\nu+a,1)^T, & \text{if } \mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{Y}}. \end{cases}$$

In particular, we now consider two special cases.

3.1 Homoscedastic case

Let us assume that $\mu^{\mathbf{Y}} \neq 0$, and $a = 1$. So, we have a common covariance for both classes. In such a case, we obtain the following:

$$\mathbf{T}_d(\mathbf{z}) \xrightarrow{P} \begin{cases} (1,\nu+1)^T, & \text{if } \mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}, \\ (\nu+1,1)^T, & \text{if } \mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{Y}}. \end{cases}$$

If $\nu > 0$, this statistic allows one to identify the distribution which produced $\mathbf{z}$. However, if $\nu = 0$, then it happens that $\mathbf{T}_d \xrightarrow{P} (1,1)^T$ independently of $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$, or $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{Y}}$. In other words, the proposed statistic allows one to decide the distribution which produced $\mathbf{z}$ without possibility of mistake, or alternatively, it is completely useless.

In [3], the authors show that the linear classifier obtains perfect classification in the case we are considering here if the series

$$\|(\Sigma^d_{\mathbf{X}})^{-1/2}\mu^{\mathbf{Y}}_d\|^2 = \sum_{i=1}^{d} (\mu^{\mathbf{Y}}_{d,i})^2/\lambda^{\mathbf{X}}_i$$

diverges as $d \to \infty$. They also prove that perfect classification is impossible if this series converges. However, if this series converges, then $(\mu^{\mathbf{Y}}_{d,i})^2/\lambda^{\mathbf{X}}_i \to 0$. In this case, it is easy to show that $\nu = 0$ and our procedure is useless. However, let us assume that $\lambda^{\mathbf{X}}_i = i^{-2}$, and $\mu^{\mathbf{Y}}_{d,i} = i^{-3/2}$. Then

$$\sum_{i\geq 1} (\mu^{\mathbf{Y}}_{d,i})^2/\lambda^{\mathbf{X}}_i = \sum_{i\geq 1} i^{-1} = \infty,$$

and the linear classifier in [3] is perfect. On the other hand, we have that

$$\frac{1}{d}\|(\Sigma^d_{\mathbf{X}})^{-1/2}\mu^{\mathbf{Y}}_d\|^2 = \frac{1}{d}\sum_{i\leq d} i^{-1} \to 0,$$

and the sequence $\{\mathbf{T}_d\}_{d\geq 1}$ is useless for classification of data points as $d \to \infty$. Therefore, if both distributions have the same covariance operator, then our procedure does not improve the linear classifier constructed by [3].

3.2 Equality in means, and difference only in scatter

We will now assume that $\mu^{\mathbf{Y}} = 0$ and $a \neq 1$. Therefore, $\mu^{\mathbf{X}} = \mu^{\mathbf{Y}} = 0$, and this implies $\nu = 0$. In this setting, we obtain the following limits

$$\mathbf{T}_d(\mathbf{z}) \xrightarrow{P} \begin{cases} (1, a^{-1})^T, & \text{if } \mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}, \\ (a, 1)^T, & \text{if } \mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{Y}}. \end{cases}$$

This is the simplest case in which two different covariance operators give way to a perfect classification problem, but one may easily construct more involved situations.

4 Classification with Unknown Distributions

In this section, we will discuss steps to implement the procedure described in Section 2. In practice, the distributions $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$, and all the involved quantities need to be estimated from the training data from the two classes.

4.1 Main result

We will use the same notation as in Section 2. For the results in this section, we will use the subspace generated by the first d eigenfunctions of the covariance operator of the random process $\mathbf{Z}$, and it will be denoted by $V_d^{\mathbf{Z}}$.

Given $s, t \in [0,1]$, we denote $K_{\mathbf{Z}}(s,t) = \mathrm{Cov}(\mathbf{Z}(s), \mathbf{Z}(t))$. Note that $K_{\mathbf{Z}}(\cdot,\cdot)$ defines a linear operator in $\mathbb{H}$ as follows:

$$(K_{\mathbf{Z}} \circ f)(t) = \int_0^1 K_{\mathbf{Z}}(s,t) f(s) ds, \text{ with } f \in \mathbb{H} \text{ and } t \in [0,1].$$

Since $K_{\mathbf{Z}}$ is a linear operator, we consider its eigenfunctions and eigenvalues, and denote them by $\phi_i^{\mathbf{Z}}(t)$ and $\lambda_i^{\mathbf{Z}}$, respectively, for $i = 1, \ldots$ It is well known that $\phi_1^{\mathbf{Z}}(t), \phi_2^{\mathbf{Z}}(t), \ldots$ forms an orthonormal basis of $\mathbb{H}$. We will now make the following assumptions:

A.1 $\sup_{t\in[0,1]} E[(\mathbf{Z}(t))^4] < \infty$.

A.2 We have $\lambda_1^{\mathbf{Z}} > \lambda_2^{\mathbf{Z}} > \ldots > 0$ satisfying $\sum_i \lambda_i^{\mathbf{Z}} < \infty$.

To estimate $K_{\mathbf{Z}}$ and its eigenvalues and eigenfunctions, we will use the corresponding empirical quantities. Suppose that we have a simple random sample $\mathbf{z}_1, \ldots, \mathbf{z}_n$ taken from $\mathbb{P}_{\mathbf{Z}}$. Given $s, t \in [0,1]$, we define

$$\hat{K}_{\mathbf{Z}}(s,t) = \frac{1}{n} \sum_{i=1}^{n} [\mathbf{z}_i(s) - \overline{\mathbf{z}}_n(s)][(\mathbf{z}_i(t) - \overline{\mathbf{z}}_n(t)],$$

where $\bar{\mathbf{z}}_n(t) = \frac{1}{n}\sum_{i=1}^n \mathbf{z}_i(t)$. Consider the corresponding estimated families $\hat{\lambda}_1^{\mathbf{Z}} \geq \hat{\lambda}_2^{\mathbf{Z}} \geq \ldots$ and $\hat{\phi}_1^{\mathbf{Z}}, \hat{\phi}_2^{\mathbf{Z}}, \ldots$ of its eigenvalues and eigenvectors, respectively. Notice that $\hat{K}_{\mathbf{Z}}$ as well as all the $\lambda_i^{\mathbf{Z}}$'s and $\hat{\phi}_i^{\mathbf{Z}}$'s depend on n. Given $\mathbf{u} \in \mathbb{H}$, we have

$$\hat{u}_i^{\mathbf{Z}} = \langle \mathbf{u}, \hat{\phi}_i^{\mathbf{Z}} \rangle = \int_0^1 \mathbf{u}(t)\hat{\phi}_i^{\mathbf{Z}}(t)dt, \text{ for } i = 1, \ldots.$$

With a finite sample, we cannot estimate all the involved eigenvalues and eigenvectors with the limit on $d \to \infty$. Thus, we follow the work in [3] and [9], and select a non-random decreasing sequence $\{\eta_n^{\}}$ going to zero slowly enough as to satisfy $\lim_n n^{1/6}\eta_n = \infty$. We take

$$R_n^{\mathbf{Z}} = \inf\{j : \hat{\lambda}_j^{\mathbf{Z}} - \hat{\lambda}_{j+1}^{\mathbf{Z}} < \eta_n\} - 1. \tag{7}$$

We will now study the map $\hat{D}^{\mathbf{Z}}_{R_n^{\mathbf{Z}}}(\mathbf{u}, \mathbf{v})$, which is the function $D_d^{\mathbf{Z}}(\mathbf{u}, \mathbf{v})$ defined in (1) with $d = R_n^{\mathbf{Z}}$ and the covariance matrix $\Sigma_d^{\mathbf{Z}}$ is estimated by $\hat{\Sigma}_d^{\mathbf{Z}}$.

Theorem 4.1 *Let $\{\mathbf{x}_n\}_{n\geq 1}$ and $\{\mathbf{y}_m\}_{m\geq 1}$ be sequences of independent observations taken from $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$, respectively. Assumptions A.1 and A.2, and those in Theorem 2.1 hold. Denote $\{\bar{\mathbf{x}}_n\}_{n\geq 1}$ and $\{\bar{\mathbf{y}}_m\}_{m\geq 1}$ to be the associated empirical means. Let $\mathbf{z}$ be an observation independent from both samples. If $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$, then*

$$\hat{D}^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n) \xrightarrow{P} 1 \text{ as } n \to \infty \text{ and } \hat{D}^{\mathbf{Y}}_{R_m^{\mathbf{Y}}}(\mathbf{z}, \bar{\mathbf{y}}_m) \xrightarrow{P} (L_\mu^{\mathbf{X},\mathbf{Y}} + L_S^{\mathbf{X},\mathbf{Y}}) \text{ as } m \to \infty.$$

If $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{Y}}$, under the obvious modification of the assumptions in Theorem 2.1 we get

$$\hat{D}^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n) \xrightarrow{P} (L_\mu^{\mathbf{Y},\mathbf{X}} + L_S^{\mathbf{Y},\mathbf{X}}) \text{ as } n \to \infty \text{ and } \hat{D}^{\mathbf{Y}}_{R_m^{\mathbf{Y}}}(\mathbf{z}, \bar{\mathbf{y}}_m) \xrightarrow{P} 1 \text{ as } m \to \infty.$$

Practical interest of Theorem 4.1 is clear from the comments after Theorem 2.1. Consequently, it happens that this result provides a procedure which asymptotically, as $\inf(n, m) \to \infty$ allows one to classify observations without a possibility of mistake.

4.2 The classification procedure in practice

We describe how our procedure works in practice. Given training samples $\mathbf{x}_1, \ldots, \mathbf{x}_n$ and $\mathbf{y}_1, \ldots, \mathbf{y}_m$ from the two classes, the classification procedure is as follows:

1. Estimate the eigenvalues, the eigenfunctions and the sample means from each of the two training samples.
2. Choose $d_{n,m}$ following the restrictions in Proposition.
3. Classify $\mathbf{z}$ as produced by $\mathbb{P}_{\mathbf{X}}$ iff $\hat{\delta}_{d_{n,m}}(\mathbf{z}) = 1$. Here $\hat{\delta}_{d_{n,m}}(\mathbf{z})$ is defined as follows:

$$\hat{\delta}_{d_{n,m}}(\mathbf{z}) = \begin{cases} 1, & \text{if } |\hat{D}^{\mathbf{X}}_{d_{n,m}}(\mathbf{z}, \bar{\mathbf{x}}_n) - 1| \leq |\hat{D}^{\mathbf{Y}}_{d_{n,m}}(\mathbf{z}, \bar{\mathbf{y}}_m) - 1|, \\ 2, & \text{else.} \end{cases}$$

Obviously, it is not possible to estimate a number of eigenfunctions higher than the sample size. If we want to apply Theorem 4.1 to be sure of the consistency of the procedure, we need to guarantee that only a specific set of values of $d_{n,m}$ are used in step 2. However, it happens that we can choose any value for $d_{n,m}$ as long as it goes to infinity with $\min(n, m)$ and not too quickly as shown in the next proposition.

Proposition 4.2 *Let $\{\mathbf{x}_n\}_{n\geq 1}$ and $\{\mathbf{y}_{m_n}\}_{n\geq 1}$ be sequences of independent observations taken from $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$, respectively, with $\lim_{n\to\infty} m_n = \infty$. Let $\{\eta_n^{\mathbf{X}}\}$ and $\{\eta_{m_n}^{\mathbf{Y}}\}$ be decreasing sequences going to zero slowly enough as to satisfy $\lim_{n,m} \min(n^{1/6}\eta_n^{\mathbf{X}}, m_n^{1/6}\eta_{m_n}^{\mathbf{Y}}) = \infty$. For each $n \in \mathbb{N}$, let $R_n^{\mathbf{X}}$ and $R_{m_n}^{\mathbf{Y}}$ be defined as in (7). Let $\{u_n\}$ be a sequence of positive numbers such that $\lim_n u_n = \infty$ and $u_n \leq \min(R_n^{\mathbf{X}}, R_{m_n}^{\mathbf{Y}})$. Choose $d_n \in \{u_n, \ldots, \min(R_n^{\mathbf{X}}, R_{m_n}^{\mathbf{Y}})\}$. Under the assumptions in Theorem 4.1, the misclassification probability of the classifier $\hat{\delta}_{d_{n,m_n}}(\mathbf{z})$ goes to 0 asymptotically as $n \to \infty$.*

The proof is stated in the Appendix. The modification to deal with $J(> 2)$ classes adds no special difficulty to our procedure. When only two classes are involved, we handle two-dimensional vectors in the classification procedure. If J populations are involved, we need to consider J-dimensional vectors. Technically, the situation when $J > 2$ is identical to the case when $J = 2$. In fact, the proofs we gave for Theorems 2.1, 2.2 and 4.1, and Proposition 4.2 allows one to prove an analogous theorem for the J class classification problem as well.

4.3 On the cross-validation approach

An important point is the selection of the optimal dimension for the projected space. For large values of $d_{n,m}$, we expect the values of the estimated statistic $\hat{\mathbf{T}}_{d_{n,m}}$ to form two clearly separated clusters depeneding on the class label of the observation. Moreover, $\hat{\mathbf{T}}_{d_{n,m}}(\mathbf{z})$ will lie closer to the cluster formed by the transformed observations with the same class label as $\mathbf{z}$ (also see Theorem 4.1). We now construct a sequence of images to show how the separation varies with increasing values of $d_{n,m}$. Consider data generated with two different Brownian motions with the same covariance operators (Example II for the location case in Section 5) for varying values of $d_{n,m}$. We have samples of size 50 from each of the two classes. The combined sample of size 100 has been used to construct an estimate for the scatter. Figure 1 clearly shows that the separation between the two classes increases with $d_{n,m}$. Observe that the transformed data points are concentrated about two distinct points corresponding to the two classes till $d_{n,m} = 60$. The data clouds start to disperse when $d_{n,m} = 80$. This can be explained by the fact that we start to observe numerical instability when we increase $d_{n,m} = 80$ with a fixed sample size $n + m = 100$.

We may use the expressions related with $R_n^{\mathbf{X}}$ and $R_m^{\mathbf{Y}}$, however we prefer an approach which directly relates to the misclassification probability. So, in practice, we propose to choose a unique dimension $d_{n,m}$ for both samples using cross-validation (CV). As we observe from Figure 1, the estimation of $d_{n,m}$ is quite crucial as it determines the dimension of the space where we project our observations for a fixed sample size. We estimate it by

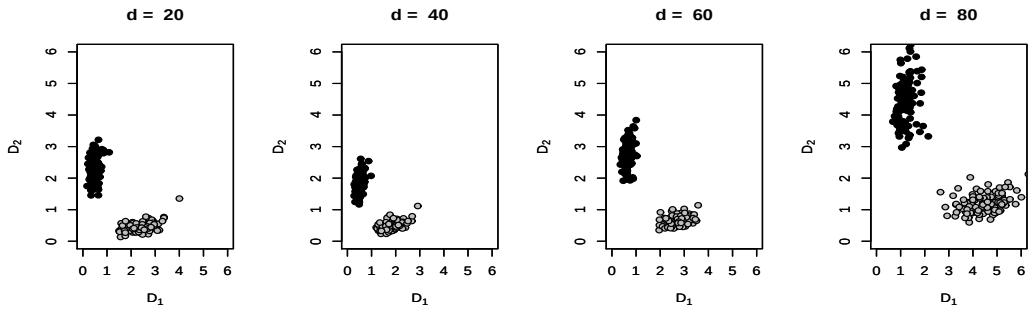

Figure 1: The grey and the black points correspond to the transformed data points of the first and the second classes, respectively.

minimizing the CV estimate [10] of the misclassification rate, and call it $\hat{d}_{CV}$ as follows. For CV, we have used a common value of d over both classes to reduce the computational burden in all situations. To estimate the misclassification rate that a fixed value of d produces, we randomly create a partition of the training sample into two subsets, S_b and its complement. The size of S_b is 0.9 times the size of the initial training sample, and the subsamples from each distribution are proportional to the original sample sizes (see [10] for more details). For each value of $b = 1, \ldots, B$, we treat the points in S_b as the training set, and the complement of the set S_b as the test set. We use S_b to obtain an estimate of the misclassification rate, repeat this partition $B = 50$ times and average it over these B samples to get $\hat{p}_d^{CV}$. Define $\hat{d}_{CV} = \arg\min_{2 \le d \le N} \hat{p}_d^{CV}$, where $N = n + m$ or $N = \min\{n, m\}$ if we use a pooled or individual estimates of the covariance, respectively.

5 Numerical Results

For our simulation study, we consider only two class problems. The training sample size of each class is 50, while the test is 100. We replicate our experiment 200 times, and the results are reported in tables below. The minimum misclassification rate is marked in **bold**, while the second lowest is in *italics*. We generated data on a discrete grid of 100 equi-spaced points in the unit interval $[0, 1]$ from four different simulation examples described below.

I. The first example is taken from [3]. The authors considered a stochastic representation $\mathbf{Z}(t) = \sum_{j=1}^{40} (\lambda_j^{1/2} Z_j + \mu_j^{\mathbf{Z}}) \phi_j(t)$. Here, the Z_js are independent standard normal random variables, $\phi_j(t) = \sqrt{2}\sin(\pi j t)$, with $t \in [0, 1]$ and $\lambda_j = 1/j^2$ for $j = 1, \ldots, 40$.

II. Let $\mathbf{Z}(t)$ with $t \in [0, 1]$ be the shifted Brownian motion having mean function $\mu^{\mathbf{Z}}(t)$ and co-variance $\Sigma^{\mathbf{Z}}(s, t) = \min\{s, t\}$ with $0 \le s, t \le 1$. Data for this example was simulated using the function `rmvnorm` from the R package `mvtnorm` [7].

III. As a third example, we generated data from the Ornstein-Uhlenbeck or Vasicek process stationary law using the linear stochastic differential equation: $d\mathbf{Z}(t) = (\theta_1 - \theta_2\mathbf{Z}(t))dt + \theta_3 d\mathbf{W}(t)$, where $\theta_2 > 0$ (ensures stationarity) and $\theta_3 > 0$ and $\mathbf{W}(t)$ is a standard Brownian motion. Denote $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)^T$. We simulated data for this GP using algorithms implemented in the function `rsOU` from the R package `sde` [11].

IV. Another interesting example was considered by [3] where the authors took the same representation as in Example I but with $\lambda_j = \exp[-\{2.1 - (j-1)/20\}^2]$ for $j = 1, \ldots, 40$. Here, it happens that $\lambda_1 < \lambda_2 < \ldots < \lambda_{40}$, and unlike assumption A.2 high valued eigenvalues occur at the higher indices.

In addition to the classifier $\hat{\delta}_{d_{n,m}}$, we may consider any classifier which separates the two clusters (see Figure 1). We consider two *linear classifiers*, namely, the centroid based classifier (CND) and support vector machines (SVM) with a linear kernel on the transformed data. One may refer to the book [10] for more details on these linear classifiers.

We consider two broad methods for comparison. The first method is based on the centroid based classifiers developed in [3]. The authors considered two classifiers, the first is based on principal components (we call it DH-PC), while the second is based on partial least squares (we call it DH-PLS). The second set of methods are non-parametric approaches developed in [4]. The authors developed two classifiers, which are available as the functions `funopare.knn.gcv` and `funopadi.knn.lcv` in [4]. The first classifier (we call it NP1) performs functional prediction (or, regression) of a scalar response from a sample of curves using the functional kernel estimator. A global bandwidth (i.e., number of neighbors) is selected by a CV procedure. While the second classifier (we call it NP2) performs a functional discrimination of a sample of curves when a categorical response is observed (i.e., supervised classification). A local bandwidth (i.e., number of neighbours depending on the curve where the estimator is evaluated) is selected by a CV procedure. Please see the manual available at http://www.math.univ-toulouse.fr/∼ferraty/SOFTWARES/NPFDA/npfda-reference-manual.pdf for details.

5.1 Difference in locations

For Example I we set the mean function $\mu_j^{\mathbf{Z}} = 0$ for $j > 6$, and the first six components to be $(0, -0.5, 1, -0.5, 1, -0.5)^T$ and $(0, -0.75, 0.75, -0.15, 1.4, 0.1)^T$, respectively, for the two distributions $\mathbb{P}_{\mathbf{X}}$ and $\mathbb{P}_{\mathbf{Y}}$ (see [3]). The mean functions for the two classes in Example II are the constant functions 0 and 1, respectively. We set the parameters to be $\boldsymbol{\theta}_1 = (0, 1, \sqrt{2})^T$ and $\boldsymbol{\theta}_2 = (1, 1, \sqrt{2})^T$ for Example III. In Example IV, [3] considered the constant function 0 for class-1, while they set the first three components to be $(0.75, -0.75, 0.75)^T$, and $\mu_j = 0$ for $j > 3$ for class-2. Interest of the last example lies in the fact that we have difference in means among the first few components, while the largest eigenvalues are in the later components. Since the two classes have the same covariance structure, we have used a common estimator for the covariance operator in both classes.

Table 1: Misclassification rates for different GPs with standard error in brackets

GP ↓	$\hat{\delta}$	**T**-CND	**T**-SVM	DH-PC	DH-PLS	NP1	NP2
I	0.0506	*0.0198*	0.0272	0.2322	**0.0008**	0.0827	0.1523
	(0.0020)	(0.0009)	(0.0012)	(0.0055)	(0.0001)	(0.0018)	(0.0016)
II	0.0039	**0.0011**	*0.0015*	0.0229	0.1990	0.2049	0.2748
	(0.0004)	(0.0001)	(0.0002)	(0.0012)	(0.0019)	(0.0041)	(0.0012)
III	0.0164	*0.0014*	0.0021	0.1058	**0.0001**	0.1996	0.2726
	(0.0016)	(0.0002)	(0.0002)	(0.0051)	(0.0001)	(0.0039)	(0.0012)
IV	0.3016	0.2005	*0.1905*	0.3618	**0.0795**	0.1952	0.2697
	(0.0041)	(0.0033)	(0.0032)	(0.0045)	(0.0014)	(0.0038)	(0.0012)

From the numerical figures in Table 1, it is clear that our classifiers generally improve over the method DH-PC. On the other hand, DH-PLS clearly outperforms all the classifiers except in Example II. Observe that PC based methods perform better than PLS in Example II, and our method has the minimum misclassification rate here. From this analysis, it turns out that DH-PLS is quite successful in estimating the optimal subspace. Since the classes differ only in their locations, we implemented the PLS algorithm for our proposed transformation as well. The results based on the classifier $\hat{\delta}^{PLS}$ were obtained to be 0.0006 (0.0001), 0.1977 (0.0020), 0.0000 (0.0000) and 0.0817 (0.0015), respectively. We observe a significant improvement in the misclassification rates in Examples I, III and IV. These numerical results indicate that the PLS algorithm is generally quite effective in problems with difference only in their locations.

Note that the performance of DH-PC is significantly better than DH-PLS in Example II. Interestingly, DH-PC (the *Bayes' classifier*) gives a higher misclassification rate than the classifier based on the proposed transformation in Example II. Figure 2 below demonstrates this point. We have used the combined sample of size 100 to estimate the covariance operator. The left hand side graph in this figure shows points in the test sample after using the proposed transformation. The right plot shows the projection based on DH-PC proposed in [3]. This graph also includes the separating hyperplane. Clearly, samples from the two classes are completely separable in the left hand side graph, while the one on the right shows some misclassified observations. Additionally, we plot the histogram of estimated values of $\hat{d}_{CV}$ over 200 random iterations for both methods (see below the main plots). The plot on the left side is negatively skewed towards the lower values of d, while the one on the right is positively skewed. The performance of DH-PC may improve by considering a larger set of possible values for d.

5.2 Difference in location and scale

The mean functions mentioned in Section 5.1, and the distribution of the first class are kept unchanged. We introduce appropriate changes in the parameters to obtain $\Sigma^{\mathbf{Y}} = 5\Sigma^{\mathbf{X}}$. This is straight forward in Example II. In Examples I and IV, we define $\lambda_j^{\mathbf{Y}} = 5\lambda_j^{\mathbf{X}}$ for all

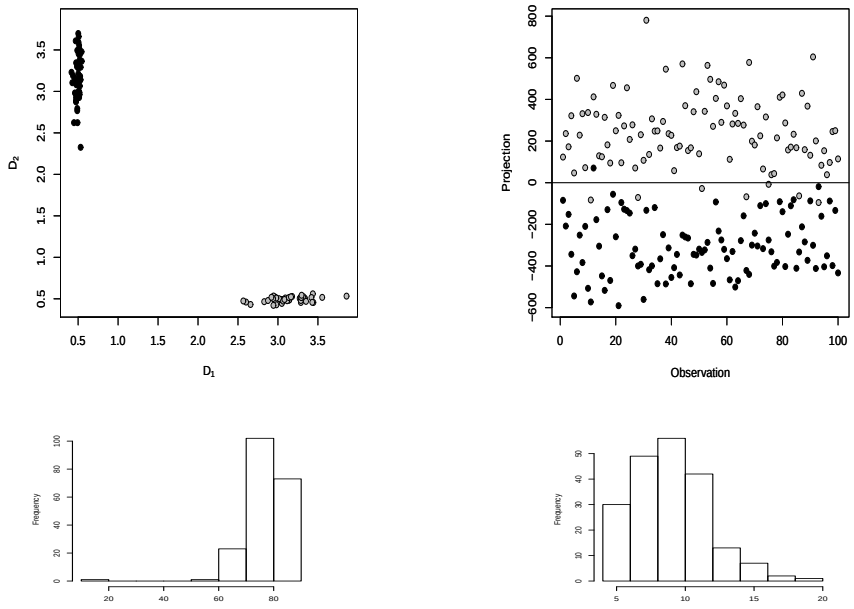

Figure 2: The data are generated from Example II. The left plot shows the two-dimensional projection of the test data points after using the transformation $\mathbf{T}$ with $\hat{d}_{CV} = 84$ (the value which minimizes $\hat{p}_d^{CV}$), while the right plot shows the projection by DH-PC and the separating line. The lower panel gives the histogram of the $\hat{d}_{CV}$ values over the 200 random iterations.

$j = 1, 2, \ldots$. The vector $\boldsymbol{\theta}_2$ is defined as $(1, 1, 5\sqrt{2})^T$ in Example III.

Table 2: Misclassification rates for different GPs with standard error in brackets

GP $\downarrow$	$\hat{\delta}$	**T**-CND	**T**-SVM	DH-PC	DH-PLS	NP1	NP2
I	0.0511	*0.0224*	**0.0133**	0.3071	0.0687	0.0556	0.1722
	(0.0012)	(0.0009)	(0.0009)	(0.0046)	(0.0014)	(0.0013)	(0.0016)
II	0.0286	*0.0188*	**0.0019**	0.1205	0.2715	0.0588	0.2155
	(0.0020)	(0.0009)	(0.0003)	(0.0025)	(0.0023)	(0.0014)	(0.0015)
III	**0.0000**	0.0082	*0.0013*	0.2715	0.1676	0.0522	0.1155
	(0.0000)	(0.0005)	(0.0002)	(0.0038)	(0.0022)	(0.0012)	(0.0011)
IV	0.1791	*0.0298*	**0.0174**	0.4005	0.2368	0.0583	0.2130
	(0.0021)	(0.0010)	(0.0011)	(0.0039)	(0.0024)	(0.0013)	(0.0015)

We now have difference in covariance structures for these examples. So, we used different covariance estimates for each of the two classes. Although DH-PLS improves over DH-PC in the examples considered here, excepting Example II, our proposed transformation brings down the misclassification rate to about 0 for all the methods. Superiority of the transformations based on $\mathbf{T}$ is quite clear from the numerical figures in Table 2.

The non-parametric classifier NP1 yields improved performance in several situations and leads to a misclassification rate of about 5%, but it is still significantly higher than those obtained using any of our methods.

5.3 Difference in scales

Here, we set the location function of the distribution corresponding to class 2 to be the constant function 0 for Examples I to IV, and retain the covariance structure mentioned in Section 5.2 for the distribution of both classes. The distributions for both classes now differ only in their scales.

Table 3: Misclassification rates for different GPs with standard error in brackets

GP ↓	$\hat{\delta}$	**T**-CND	**T**-SVM	DH-PC	DH-PLS	NP1	NP2
I	0.0811	*0.0240*	**0.0074**	0.4723	0.4172	0.0601	0.2128
	(0.0016)	(0.0009)	(0.0005)	(0.0026)	(0.0022)	(0.0014)	(0.0016)
II	*0.0052*	0.0228	**0.0028**	0.4700	0.4783	0.0629	0.2162
	(0.0004)	(0.0009)	(0.0003)	(0.0029)	(0.0027)	(0.0015)	(0.0014)
III	**0.0000**	0.0078	*0.0008*	0.3897	0.3720	0.0503	0.1142
	(0.0000)	(0.0005)	(0.0001)	(0.0033)	(0.0023)	(0.0012)	(0.0015)
IV	0.1937	*0.0294*	**0.0168**	0.4704	0.4152	0.0600	0.2120
	(0.0020)	(0.0009)	(0.0010)	(0.0027)	(0.0022)	(0.0014)	(0.0015)

The strength of our transformations is further evident from the analysis of the scale problems. The misclassification rate of all our classifiers in all four examples is again close to 0 (see Table 3), and demonstrates the '*perfect classification property*' of the proposed transformations. However, the centroid based classifier DH-PC now yields a misclassification rate close to the random tosses of an unbiased coin in several examples. This is expected in view of the fact that there is no information in locations any more. The misclassification rates of NP1 are similar to those obtained in Table 2 but statistically significant compared to 0, while the results of NP2 turn out to be not too good.

6 Analysis of Benchmark Datasets

We have applied our methods to three benchmark data sets. In each set, we had a sample of $\sum_j m_j$ observations, where we have m_j data points from the j-th population for $1 \leq j \leq J$. The data were randomly split $M = 200$ times to construct a training sample of size n_j, and a test sample of size N_j from j-th population with $n_j + N_j = m_j$ for all $1 \leq j \leq J$. Define $n = \sum_j n_j$ and $N = \sum_j N_j$. We built the empirical classifiers using the training data, and used them to classify the test data points. For each classifier, the corresponding proportion of test observations that were misclassified is reported below in Table 4.

Table 4: Misclassification rates for different real datasets with standard error in brackets

Data	(n, N)	$\hat{\delta}$	**T**-CND	**T**-SVM	DH-PC	DH-PLS	NP1	NP2
Growth	(46,47)	0.1780	*0.0595*	**0.0585**	0.1403	0.2587	0.0947	0.0838
		(0.0039)	(0.0030)	(0.0025)	(0.0032)	(0.0061)	(0.0026)	(0.0015)
Wheat	(30,70)	0.0114	**0.0011**	*0.0022*	0.3048	0.2050	0.1713	0.1217
		(0.0015)	(0.0004)	(0.0005)	(0.0047)	(0.0078)	(0.0033)	(0.0022)
	(50,50)	0.0182	**0.0001**	*0.0023*	0.1395	0.2530	0.1330	0.0904
		(0.0006)	(0.0001)	(0.0004)	(0.0066)	(0.0074)	(0.0024)	(0.0014)
Synthetic	(150,450)	*0.0789*	0.1628	**0.0394**	$*$	$*$	0.0941	$*$
Control		(0.0001)	(0.0040)	(0.0010)			(0.0014)	
	(300,300)	*0.1167*	0.1177	**0.0140**	$*$	$*$	0.2825	$*$
		(0.0002)	(0.0032)	(0.0005)			(0.0017)	

$*$ We could not run for $J = 6$ classes

The first two data sets are two class problems, while the third is a six class problem. The `growth data` [4] contains the heights of 39 *boys* and 54 *girls* from age 1 to age 18, and the ages at which they were collected. The heights were measured at 31 different time points during this age span. We distributed the data as $n_1 = 19$ and $n_2 = 27$; and $N_1 = 20$ and $N_2 = 27$. The `wheat data` consists of near infra-red spectra of 100 wheat samples with known protein content, and were measured from 1100nm to 2500nm at 2 nm intervals to yield data observed at 701 points (see [3] for more details). We have 41 samples with *protein content less than* 15, and 59 samples with *protein content higher than* 15. Two different configurations for the data split are considered here. The first is $n_1 = 12$ and $n_2 = 18$; and $N_1 = 29$ and $N_2 = 41$, while the second is $n_1 = 20$ and $n_2 = 30$; and $N_1 = 21$ and $N_2 = 29$. The `synthetic control data` is obtained from the UCI machine learning repository. It contains data of six different patterns, namely, *downward trend, cyclic, normal, upward shift, upward trend* and *downward shift*. There are 100 observations in each of the classes having time series of length 60. We have considered two splits here as: $n_i = 25$ and $N_i = 75$; and $n_i = N_i = 50$ for all $1 \leq i \leq 6$. For numerical stability, we have used a common estimate for the covariance over the competing classes.

Numerical figures for the three datasets in Table 4 demonstrate a definite point in the '*perfect classification*' scenario. The results clearly exhibit that good performance can be achieved by the proposed transformation. We observe a significant improvement in the misclassification rate over existing parametric as well as nonparametric classification methods. Among the proposed classification methods, the one based on SVM with linear kernel usually turns out to be the winner by majority vote. The performance of the method based on centroids turns out to be quite stable, and usually turns out to be the second best. Note that SVM finds the best linear separator, and this explains its overall good performance.

Acknowledgements

The authors would like to include here a mention to the late Prof. Peter Hall whose talk on the paper [3] (see https://www.youtube.com/watch?v=TLoxdL2XRBo) at the Universidad de Cantabria in January 2014 sparkled the interest of the authors on the subject. The first author also acknowledges kind conversations on the subject with him.

7 Appendix: Proofs and Mathematical Details

Proposition 7.1 *Let $\{A_d\}_d$ be a sequence of $d \times d$-dimensional, invertible matrices. Denote $\alpha_1^d, \ldots, \alpha_d^d$ to be the eigenvalues of the matrix $S_d = (A_d)^{-1/2}\Sigma_d^{\mathbf{Z}}(A_d)^{-1/2}$. Let $\mu \in \mathbb{H}$. Assume that $\mathbf{z}$ has been produced from the distribution $\mathbb{P}_{\mathbf{Z}}$, and that there exist L_μ and L_S (finite or not) such that*

$$L_\mu = \lim_{d\to\infty} \frac{1}{d}\|(A_d)^{-1/2}(\mu^{\mathbf{Z}} - \mu)_d\|^2, \tag{8}$$

$$L_S = \lim_{d\to\infty} \frac{1}{d}\ trace(S_d),\ and \tag{9}$$

$$0 = \lim_{d\to\infty} \frac{\sup(\alpha_1^d, \ldots, \alpha_d^d)}{d}. \tag{10}$$

Then

$$\frac{1}{d}\|(A_d)^{-1/2}(\mathbf{z} - \mu)_d\|^2 \xrightarrow{P} (L_\mu + L_S)\ as\ d \to \infty.$$

Proof: Notice that $(\mathbf{z}-\mu)_d$ is a d-dimensional normal vector, with mean equal to $(\mu^{\mathbf{Z}}-\mu)_d$ and covariance matrix equal to $\Sigma_d^{\mathbf{Z}}$. Thus, $\|(A_d)^{-1/2}(\mathbf{z} - \mu)_d\|^2$ is equal to the square of the norm of a d-dimensional normal variable with mean $\mathbf{m}_d = (A_d)^{-1/2}(\mu^{\mathbf{Z}} - \mu)_d$ and covariance matrix S_d. Therefore, if $\mathbf{u}_d$ is a d-dimensional vector with centered normal distribution and covariance matrix equal to S_d, then $\frac{1}{d}\|(A_d)^{-1/2}(\mathbf{z} - \mu)_d\|^2$ is equally distributed with

$$\frac{1}{d}\langle \mathbf{m}_d + \mathbf{u}_d, \mathbf{m}_d + \mathbf{u}_d\rangle = \frac{1}{d}\left(\|\mathbf{m}_d\|^2 + \|\mathbf{u}_d\|^2 + 2\langle \mathbf{m}_d, \mathbf{u}_d\rangle\right). \tag{11}$$

By assumption (8), we have

$$\lim_{d\to\infty} \frac{1}{d}\|\mathbf{m}_d\|^2 = L_\mu.$$

Let us consider the second term in (11). Fix a basis in V_d spanned by the eigenvectors of S_d. Denote $\mathbf{u}_d = (u_{d,1}, \ldots, u_{d,d})^T$ and $\mathbf{m}_d = (m_{d,1}, \ldots, m_{d,d})^T$ in this basis. Therefore, the random variables $(u_{d,i})^2$ with $1 \le i \le d$ are independent with means equal to α_i^d for $1 \le i \le d$ and $\sum_{i=1}^d (u_{d,i})^2$ has the same distribution as $\sum_{i=1}^d \alpha_i^d (u_i)^2$. Here, $\{u_i\}_{1\le i\le d}$ is a sequence of i.i.d. real variables with the standard normal distribution.

Fix $\epsilon > 0$, and assume first that L_S is finite. Taking into account that the variance of a chi-squared distribution with one degree of freedom is two and using Tchebychev's inequality, we have that

$$\begin{aligned}
\mathbb{P}\left[\frac{1}{d}\left|\|\mathbf{u}_d\|^2 - \text{trace}(S_d)\right| \geq \epsilon\right] &= \mathbb{P}\left[\frac{1}{d}\left|\sum_{i=1}^{d}\left((u_{d,i})^2 - \alpha_i^d\right)\right| \geq \epsilon\right] \\
&\leq \frac{2}{\epsilon^2 d^2}\sum_{i=1}^{d}(\alpha_i^d)^2 \\
&\leq \frac{2}{\epsilon^2 d^2}\sup(\alpha_1,\ldots,\alpha_d)\sum_{i=1}^{d}\alpha_i^d,
\end{aligned}$$

which converges to zero by assumptions (9) and (10). Consequently, we have shown that

$$\frac{1}{d}\|\mathbf{u}_d\|^2 - \frac{1}{d}\text{trace}(S_d) \xrightarrow{P} 0 \text{ as } d \to \infty,$$

and assumption (9) gives

$$\frac{1}{d}\|\mathbf{u}_d\|^2 \xrightarrow{P} L_S \text{ as } d \to \infty.$$

Assume now that L_S is infinite. We have that

$$\begin{aligned}
\mathbb{P}\left[\frac{1}{\sum_{i=1}^{d}\alpha_i^d}\left|\sum_{i=1}^{d}\left((u_{d,i})^2 - \alpha_i^d\right)\right| \geq \epsilon\right] &= \mathbb{P}\left[\left|\sum_{i=1}^{d}\frac{\alpha_i^d}{\sum_{i=1}^{d}\alpha_i^d}\left((u_i)^2 - 1\right)\right| \geq \epsilon\right] \\
&\leq \frac{2}{\epsilon^2}\sum_{i=1}^{d}\left(\frac{\alpha_i^d}{\sum_{i=1}^{d}\alpha_i^d}\right)^2 \\
&\leq \frac{2}{\epsilon^2}\frac{\sup(\alpha_1,\ldots,\alpha_d)}{\sum_{i=1}^{d}\alpha_i^d},
\end{aligned}$$

which converges to zero as $L_S = \infty$ and using assumption (10). Therefore, we have shown that

$$\frac{1}{\frac{1}{d}\sum_{i=1}^{d}\alpha_i^d}\left(\frac{1}{d}\|\mathbf{u}_d\|^2 - \frac{1}{d}\text{ trace}(S_d)\right) \xrightarrow{P} 0. \tag{12}$$

Consequently, $\frac{1}{d}\|\mathbf{u}_d\|^2$ converges to ∞ at the same rate as $\frac{1}{d}$ trace(S_d).

Concerning the last term in (11), we have $\langle \mathbf{m}_d, \mathbf{u}_d\rangle = \sum_{i=1}^{d} m_{d,i}u_{d,i}$. Let us assume that L_μ and L_S are finite. Let $\epsilon > 0$. Using Tchebychev's inequality again, we get

$$\mathbb{P}\left[\frac{1}{d}|\langle \mathbf{m}_d, \mathbf{u}_d\rangle| > \epsilon\right] \leq \frac{1}{\epsilon^2 d^2}\sum_{i=1}^{d}(m_{d,i})^2\alpha_i^d \leq \frac{1}{\epsilon^2 d^2}\sup(\alpha_1^d,\ldots,\alpha_d^d)\|\mathbf{m}_d\|^2,$$

which converges to zero by assumptions (8) and (10), and the proposition is proved in this case.

When L_μ or L_S are infinite, the result follows from (11) and the previous results, if we are able to show that the sequence of real valued random variables

$$w_d = \frac{\langle \mathbf{m}_d, \mathbf{u}_d \rangle}{\max(\|\mathbf{m}_d\|^2, \|\mathbf{u}_d\|^2)},\ d = 1, \ldots$$

converges to zero in probability as $d \to \infty$. In turn, this will be fixed if we show that every subsequence of $\{w_d\}$ contains a new subsequence which satisfies this property. Thus, let $\{w_{d_k}\}$ be a subsequence of $\{w_d\}$ and let us consider the associated subsequences $\{\|\mathbf{m}_{d_k}\|\}$ and $\{\|\mathbf{u}_{d_k}\|\}$. Obviously, there exists a further subsequence $\{d_{k^*}\}$ such that one of the following holds:

(i) $\lim_{d_{k^*}} \dfrac{\|\mathbf{m}_{d_{k^*}}\|^2}{\mathrm{trace}(S_{d_{k^*}})} = 0.$

(ii) $\lim_{d_{k^*}} \dfrac{\|\mathbf{m}_{d_k}\|^2}{\mathrm{trace}(S_{d_{k^*}})} = \infty.$

(iii) There exists a finite $C > 0$ such that $\lim_{d_{k^*}} \dfrac{\|\mathbf{m}_{d_{k^*}}\|^2}{\mathrm{trace}(S_{d_{k^*}})} = C.$

To simplify notation, we denote the sequence $\{S_{d_{k^*}}\}$ by $\{S_h\}$, and similarly for the remaining ones. In case (i), L_S must be infinite, since (12) shows that

$$\frac{\|\mathbf{u}_h\|^2}{\mathrm{trace}(S_h)} \xrightarrow{P} 1 \text{ as } h \to \infty, \tag{13}$$

we have $\dfrac{\|\mathbf{m}_h\|}{\|\mathbf{u}_h\|} \xrightarrow{P} 0$ as $h \to \infty$. Consequently,

$$\lim_h |w_h| = \lim_h \frac{|\langle \mathbf{m}_h, \mathbf{u}_h \rangle|}{\|\mathbf{u}_h\|^2} \leq \lim_h \frac{\|\mathbf{m}_h\|}{\|\mathbf{u}_h\|} = 0 \text{ in probability.}$$

If (ii) holds, we have that $|w_h| \leq \frac{\|\mathbf{u}_h\|}{\|\mathbf{m}_h\|}$. Since $E\left[\|\mathbf{u}_h\|^2\right] = \mathrm{trace}(S_d)$, we have that $\frac{\|\mathbf{u}_h\|^2}{\|\mathbf{m}_h\|^2} \to 0$ in probability, and, also in this case $w_h \xrightarrow{P} 0$ as $h \to \infty$.

In case (iii), we also have that L_S to be infinite. Thus, taking into account that (13) also holds now, it happens that we need to show that

$$\frac{\langle \mathbf{m}_h, \mathbf{u}_h \rangle}{(C - \delta)\ \mathrm{trace}(S_h)} \xrightarrow{P} 0 \text{ as } h \to \infty,$$

where δ is an arbitrary number satisfying $0 < \delta < C$. Fix $\epsilon > 0$. We have that $\langle \mathbf{m}_h, \mathbf{u}_h \rangle$ has the same distribution as

$$\sum_{i=1}^{h} (\alpha_i^h)^{1/2} m_{h,i} u_i,$$

and consequently, we have that

$$\begin{aligned}\mathbb{P}\left[\left|\frac{\langle \mathbf{m}_h, \mathbf{u}_h\rangle}{(C-\delta)\ \mathrm{trace}(S_h)}\right| > \epsilon\right] &\leq \frac{1}{(C-\delta)^2\epsilon^2}\sum_{i=1}^{h}\frac{m_{h,i}^2\alpha_i^h}{\left(\sum_{i=1}^h \alpha_i^h\right)^2}\\ &\leq \frac{1}{(C-\delta)^2\epsilon^2}\frac{\max(\alpha_1^h,\ldots,\alpha_h^h)}{\sum_{i=1}^h\alpha_i^h}\sum_{i=1}^{h}\frac{m_{h,i}^2}{\sum_{i=1}^h\alpha_i^h}\\ &= \frac{1}{(C-\delta)^2\epsilon^2}\frac{\max(\alpha_1^h,\ldots,\alpha_h^h)}{\sum_{i=1}^h\alpha_i^h}\frac{\|\mathbf{m}_h\|^2}{\mathrm{trace}(S_h)},\end{aligned}$$

which converges to zero by assumptions (9) and (10). •

Lemma 7.2 *Let $\{a_d\}_{d\geq 1}$ be a sequence of real positive numbers such that $\lim_d \frac{1}{d}\sum_{i=1}^d a_i$ exists, and it is finite. Then, it happens that*

$$0 = \lim_d \frac{1}{d}\max_{i=1,\ldots,d} a_i.$$

Proof: Let d be a natural number, and denote $A_d = \sum_{i=1}^d a_i$. We have that

$$\frac{a_d}{d} = \frac{A_d}{d} - \frac{A_{d-1}}{d-1}\frac{d-1}{d}$$

and consequently, $0 = \lim_d \frac{a_d}{d}$. Given $\epsilon > 0$, there exists $d_1 > 0$ such that if $d > d_1$, then $\frac{a_d}{d} \leq \epsilon$ and $d_2 \geq d_1$ such that

$$\sup_{i=1,\ldots,d_1}\frac{a_i}{d_2} \leq \epsilon.$$

Let $d > d_2$ and take $1 \leq i \leq d$. So, we have that if $i \leq d_1$, then $\frac{a_i}{d} < \frac{a_i}{d_2} \leq \epsilon$ and if $i > d_1$, then $\frac{a_i}{d} \leq \frac{a_i}{i} \leq \epsilon$. This completes the proof. •

7.1 Proof of Theorem 4.1

Recall that in this theorem, we use the subspaces generated by the estimates of the first d eigenfunctions of the covariance operator of the random process $\mathbf{Z}$.

We begin with some notation and preliminary results which have been taken from [3] and [9], or follow directly from the results there. Then, we will give the proof of Theorem 4.1. For every $n \in \mathbb{N}$, let us consider

$$\begin{aligned}\hat{\Delta}_{\mathbf{Z}}^2 &= \int_0^1\int_0^1 (\hat{K}_{\mathbf{Z}}(s,t) - K_{\mathbf{Z}}(s,t))^2 dsdt,\\ \delta_j^{\mathbf{Z}} &= \min_{k\leq j}(\lambda_k^{\mathbf{Z}} - \lambda_{k+1}^{\mathbf{Z}}).\end{aligned}$$

In [3] and [9], it is shown that if $j \leq R_n^{\mathbf{Z}}$ (recall the definition of $R_n^{\mathbf{Z}}$ in (7)), then

$$\begin{aligned} |\hat{\lambda}_j^{\mathbf{Z}} - \lambda_j^{\mathbf{Z}}| &\leq \hat{\Delta}_{\mathbf{Z}}, & (14)\\ \|\hat{\phi}_j^{\mathbf{Z}} - \phi_j^{\mathbf{Z}}\| &\leq 8^{1/2}\hat{\Delta}_{\mathbf{Z}}(\delta_j^{\mathbf{Z}})^{-1}, & (15)\\ \hat{\Delta}_{\mathbf{Z}} &= O_p(n^{-1/2}), & (16)\\ R_n^{\mathbf{Z}} \to \infty \text{ and } R_n^{\mathbf{Z}} &\leq \hat{\lambda}_1^{\mathbf{Z}}\eta_n^{-1}. & (17) \end{aligned}$$

Moreover, if $j \leq R_n^{\mathbf{Z}}$, there exists a $k \leq j$ such that

$$\delta_j^{\mathbf{Z}} = \lambda_k^{\mathbf{Z}} - \lambda_{k+1}^{\mathbf{Z}} \geq \hat{\lambda}_k^{\mathbf{Z}} - \hat{\lambda}_{k+1}^{\mathbf{Z}} - 2\hat{\Delta}_{\mathbf{Z}} \geq \eta_n - 2\hat{\Delta}_{\mathbf{Z}} = \eta_n + o_P(\eta_n), \quad (18)$$

where we have applied (14) and (7) and that, from (16) and the assumption on η_n, we can conclude that $\eta_n > 2\hat{\Delta}_{\mathbf{Z}}$ from an index onward. Thus, (18) and (15) yield

$$\|\hat{\phi}_j^{\mathbf{Z}} - \phi_j^{\mathbf{Z}}\| \leq 8^{1/2}\frac{\hat{\Delta}_{\mathbf{Z}}}{\eta_n - 2\hat{\Delta}_{\mathbf{Z}}}. \quad (19)$$

From (14), (7) and (16), we obtain that

$$\lambda_j^{\mathbf{Z}} \geq \hat{\lambda}_j^{\mathbf{Z}} - \hat{\Delta}_{\mathbf{Z}} \geq \eta_n - \hat{\Delta}_{\mathbf{Z}} = \eta_n + o_P(\eta_n). \quad (20)$$

Now, we are in a position to prove Theorem 4.1.

Proof of Theorem 4.1: Let us assume that $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$. Without loss of generality, we assume that $\mu^{\mathbf{X}} = 0$. We will split the proof into two lemmas. In Lemma 7.3, we will first prove that $\left|\hat{D}^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n) - D^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n)\right| \xrightarrow{P} 0$ as $n \to \infty$. The proof that $\left|\hat{D}^{\mathbf{Y}}_{R_m^{\mathbf{Y}}}(\mathbf{z}, \bar{\mathbf{y}}_m) - D^{\mathbf{Y}}_{R_m^{\mathbf{Y}}}(\mathbf{z}, \bar{\mathbf{y}}_m)\right| \xrightarrow{P} 0$ as $m \to \infty$ is identical.

Then, we will show in Lemma 7.4 that the limits of $D^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n)$ and $D^{\mathbf{Y}}_{R_m^{\mathbf{Y}}}(\mathbf{z}, \bar{\mathbf{y}}_n)$ coincide with those of $D^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \mu^{\mathbf{X}})$ and $D^{\mathbf{Y}}_{R_m^{\mathbf{Y}}}(\mathbf{z}, \mu^{\mathbf{Y}})$, respectively.

Lemma 7.3 *Under the assumptions in Theorem 4.1, it happens that*

$$\left|\hat{D}^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n) - D^{\mathbf{X}}_{R_n^{\mathbf{X}}}(\mathbf{z}, \bar{\mathbf{x}}_n)\right| \xrightarrow{P} 0 \text{ as } n \to \infty.$$

Proof. For a fixed $\mathbf{z}$, let us denote $\mathbf{u} = \mathbf{z} - \bar{\mathbf{x}}_n$. Note that $\mathbf{u}$ depends on n, but by the Strong Law of Large Numbers, we have that $\|\mathbf{u}\| \leq \|\mathbf{z}\| + \|\bar{\mathbf{x}}_n\| = O_P(1)$. Let $n \in \mathbb{N}$, and take $j \leq R_n^{\mathbf{X}}$. We have that

$$\begin{aligned} \left|\frac{\left(u_j^{\mathbf{X}}\right)^2}{\lambda_j^{\mathbf{X}}} - \frac{\left(\hat{u}_j^{\mathbf{X}}\right)^2}{\hat{\lambda}_j^{\mathbf{X}}}\right| &= \left|\frac{u_j^{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}} - \frac{\hat{u}_j^{\mathbf{X}}}{(\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right|\left|\frac{u_j^{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}} + \frac{\hat{u}_j^{\mathbf{X}}}{(\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right| \\ &\leq \left(\left|\frac{u_j^{\mathbf{X}} - \hat{u}_j^{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}}\right| + \left|\hat{u}_j^{\mathbf{X}}\frac{(\lambda_j^{\mathbf{X}})^{1/2} - (\hat{\lambda}_j^{\mathbf{X}})^{1/2}}{(\lambda_j^{\mathbf{X}}\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right|\right)\left|\frac{u_j^{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}} + \frac{\hat{u}_j^{\mathbf{X}}}{(\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right|. \end{aligned}$$

We analyze each term in this expression separately as follows:

$$\begin{aligned}\left|\frac{u_j^{\mathbf{X}}-\hat{u}_j^{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}}\right| &\leq \frac{1}{(\lambda_j^{\mathbf{X}})^{1/2}}\int_0^1 |\mathbf{u}(t)||\phi_j^{\mathbf{X}}(t)-\hat{\phi}_j^{\mathbf{X}}(t)|dt \\ &\leq \frac{\|\mathbf{u}\|\ \|\phi_j^{\mathbf{X}}-\hat{\phi}_j^{\mathbf{X}}\|}{(\lambda_j^{\mathbf{X}})^{1/2}} \\ &\leq 8^{1/2}\|\mathbf{u}\|\frac{\hat{\Delta}_{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}(\eta_n-2\hat{\Delta}_{\mathbf{X}})} \\ &\leq 8^{1/2}\|\mathbf{u}\|\hat{\Delta}_{\mathbf{X}}(\eta_n^{-3/2}+o_P(\eta_n^{-3/2})), \end{aligned} \tag{21}$$

where we have applied the Cauchy-Schwartz inequality, (19), (16) and (20). On the other hand, we have

$$\begin{aligned}\left|\hat{u}_j^{\mathbf{X}}\frac{(\lambda_j^{\mathbf{X}})^{1/2}-(\hat{\lambda}_j^{\mathbf{X}})^{1/2}}{(\lambda_j^{\mathbf{X}}\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right| &\leq \int |\mathbf{u}(t)||\hat{\phi}_j^{\mathbf{X}}(t)|dt\frac{|\lambda_j^{\mathbf{X}}-\hat{\lambda}_j^{\mathbf{X}}|}{\left((\lambda_j^{\mathbf{X}})^{1/2}+(\hat{\lambda}_j^{\mathbf{X}})^{1/2}\right)(\lambda_j^{\mathbf{X}}\hat{\lambda}_j^{\mathbf{X}})^{1/2}} \\ &\leq \|\mathbf{u}\|\frac{\hat{\Delta}_{\mathbf{X}}}{\left((\lambda_j^{\mathbf{X}})^{1/2}+(\hat{\lambda}_j^{\mathbf{X}})^{1/2}\right)(\lambda_j^{\mathbf{X}}\hat{\lambda}_j^{\mathbf{X}})^{1/2}} \\ &\leq \|\mathbf{u}\|\hat{\Delta}_{\mathbf{X}}(\eta_n^{-3/2}+o_P(\eta_n^{-3/2})), \end{aligned} \tag{22}$$

where we have applies (7) and (20). Concerning the final term, using (20) and (7) again, we obtain that

$$\left|\frac{u_j^{\mathbf{X}}}{(\lambda_j^{\mathbf{X}})^{1/2}}+\frac{\hat{u}_j^{\mathbf{X}}}{(\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right| \leq \|\mathbf{u}\|\left(\frac{1}{(\lambda_j^{\mathbf{X}})^{1/2}}+\frac{1}{(\hat{\lambda}_j^{\mathbf{X}})^{1/2}}\right) \leq \|\mathbf{u}\|(\eta_n^{-1/2}+o_P(\eta_n^{-1/2})). \tag{23}$$

Now, if we define $C=8^{1/2}+1$, combining (21), (22), (23), (17) and (16), we get the following:

$$\begin{aligned}\left|\hat{D}_{R_n^{\mathbf{X}}}^{\mathbf{X}}(\mathbf{z},\bar{\mathbf{x}}_n)-D_{R_n^{\mathbf{X}}}^{\mathbf{X}}(\mathbf{z},\bar{\mathbf{x}}_n)\right| &\leq \sum_{j=1}^{R_n^{\mathbf{X}}}\left|\frac{\left(u_j^{\mathbf{X}}\right)^2}{\lambda_j^{\mathbf{X}}}-\frac{\left(\hat{u}_j^{\mathbf{X}}\right)^2}{\hat{\lambda}_j^{\mathbf{X}}}\right| \\ &\leq CR_n^{\mathbf{X}}\|\mathbf{u}\|^2\hat{\Delta}_{\mathbf{X}}(\eta_n^{-2}+o_P(\eta_n^{-2})) \\ &\leq C\|\mathbf{u}\|^2\hat{\lambda}_1^{\mathbf{Z}}\hat{\Delta}_{\mathbf{X}}(\eta_n^{-3}+o_P(\eta_n^{-3}))=O_P(n^{-1/2}\eta_n^{-3}). \end{aligned}$$

By construction, η_n is such that $n\eta_n^6\to\infty$. So, we have $|\hat{D}_{R_n^{\mathbf{X}}}^{\mathbf{X}}(\mathbf{z},\bar{\mathbf{x}}_n)-D_{R_n^{\mathbf{X}}}^{\mathbf{X}}(\mathbf{z},\bar{\mathbf{x}}_n)|\xrightarrow{P}0$ as $n\to\infty$, and this lemma is proved. •

Lemma 7.4 *Under the assumptions in Theorem 4.1, it happens that the limits in probability of $D_{R_n^{\mathbf{X}}}^{\mathbf{X}}(\mathbf{z},\bar{\mathbf{x}}_n)$ and $D_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}(\mathbf{z},\bar{\mathbf{y}}_n)$ coincide with that of $D_{R_n^{\mathbf{X}}}^{\mathbf{X}}(\mathbf{z},\mu^{\mathbf{X}})$ and $D_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}(\mathbf{z},\mu^{\mathbf{Y}})$, respectively.*

Proof. We will first show that $D^{\mathbf{X}}_{R^{\mathbf{X}}_n}(\mathbf{z},\bar{\mathbf{x}}_n) - D^{\mathbf{X}}_{R^{\mathbf{X}}_n}(\mathbf{z},\mu^{\mathbf{X}}) \xrightarrow{P} 0$ as $n\to\infty$. To this, let us denote $\mathbf{z} = (z_1,\ldots,z_{R^{\mathbf{X}}_n})^T, \bar{\mathbf{x}}_n = (\bar{x}_1,\ldots,\bar{x}_{R^{\mathbf{X}}_n})^T$, when written in the basis formed by the eigenvectors of $\Sigma^{\mathbf{X}}_{R^{\mathbf{X}}_n}$. Since we are assuming that $\mu^{\mathbf{X}}=0$, we have

$$D^{\mathbf{X}}_{R^{\mathbf{X}}_n}(\mathbf{z},\bar{\mathbf{x}}_n) - D^{\mathbf{X}}_{R^{\mathbf{X}}_n}(\mathbf{z},\mu^{\mathbf{X}}) \;=\; \frac{1}{R^{\mathbf{X}}_n}\sum_{j=1}^{R^{\mathbf{X}}_n}\frac{(\bar{x}_j)^2}{\lambda^{\mathbf{X}}_j} - \frac{2}{R^{\mathbf{X}}_n}\sum_{j=1}^{R^{\mathbf{X}}_n}\frac{z_j\bar{x}_j}{\lambda^{\mathbf{X}}_j}. \tag{24}$$

Since $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$, it happens that $\{z_j/(\lambda^{\mathbf{X}}_j)^{1/2}\}$ are i.i.d. with the standard normal distribution, and they are independent from the i.i.d. variables $\{\bar{x}_j(n/\lambda^{\mathbf{X}}_j)^{1/2}, 1\le j\le R^{\mathbf{X}}_n\}$ whose distribution is also standard normal. Thus, if we take two independent sequences of i.i.d. standard normal variates $\{N^1_j\}$ and $\{N^2_j\}$, it happens that

$$\frac{2}{R^{\mathbf{X}}_n}\sum_{j=1}^{R^{\mathbf{X}}_n}\frac{z_j\bar{x}_j}{\lambda^{\mathbf{X}}_j} \overset{D}{=} \frac{2}{R^{\mathbf{X}}_n(n^{1/2})}\sum_{j=1}^{R^{\mathbf{X}}_n} N^1_j N^2_j,$$

which converges in probability to zero by SLLN because according to (17), $R^{\mathbf{X}}_n\to\infty$. Thus, the second term in the left hand side of (24) converges to zero in probability. The reasoning to prove that the first term in the left hand side in (24) converges to zero in probability is similar to the previous one, taking into account that the variables $\left\{n(\bar{x}_j)^2/\lambda^{\mathbf{X}}_j, 1\le j\le R^{\mathbf{X}}_n\right\}$ are i.i.d. with χ^2 distribution with one degree of freedom.

We will now show that the limits of the sequences $\{D^{\mathbf{Y}}_{R^{\mathbf{Y}}_m}(\mathbf{z},\bar{\mathbf{y}}_m)\}$ and $\{D^{\mathbf{Y}}_{R^{\mathbf{Y}}_m}(\mathbf{z},\mu^{\mathbf{Y}})\}$ coincide in probability as $m\to\infty$. In this part, we will change the notation and we will denote $\mathbf{z} = (z_1,\ldots,z_{R^{\mathbf{Y}}_m})^T$ and $\bar{\mathbf{y}}_n = (\bar{y}_1,\ldots,\bar{y}_{R^{\mathbf{Y}}_m})^T$ when written in the basis formed by the eigenvectors of $\Sigma^{\mathbf{Y}}_{R^{\mathbf{Y}}_m}$. We begin assuming that $L^{\mathbf{X},\mathbf{Y}}_S$ and $L^{\mathbf{X},\mathbf{Y}}_\mu$ are finite.

Let us consider a non-random sequence $\{T_m\}\subset\mathbb{N}$, going to infinity with exact order $m^{2/6}$. For every $\epsilon>0$, we have that

$$\begin{aligned}
&\mathbb{P}\left[\left|D^{\mathbf{Y}}_{R^{\mathbf{Y}}_m}(\mathbf{z},\bar{\mathbf{y}}_n) - D^{\mathbf{Y}}_{R^{\mathbf{Y}}_m}(\mathbf{z},\mu^{\mathbf{Y}})\right| > \epsilon\right] \\
&\quad\le\; \mathbb{P}[R^{\mathbf{Y}}_m > T_m] + \mathbb{P}\left[\frac{1}{R^{\mathbf{Y}}_m}\sum_{i=1}^{T_m}\left|\frac{(z_j-\bar{y}_j)^2-(z_j-\mu^{\mathbf{Y}}_j)^2}{\lambda^{\mathbf{Y}}_j}\right| > \epsilon\right].
\end{aligned}$$

From (17), we have that the first term here converges to zero. Therefore, to finish this step, we only need to prove that the sequence $\left\{\frac{1}{R^{\mathbf{Y}}_m}\sum_{i=1}^{T_m}\left|\frac{(z_j-\bar{y}_j)^2-(z_j-\mu^{\mathbf{Y}}_j)^2}{\lambda^{\mathbf{Y}}_j}\right|\right\}_{m\ge1}$ converges to zero in probability. However,

$$\begin{aligned}
\sum_{j=1}^{T_m}\left|\frac{(z_j-\bar{y}_j)^2-(z_j-\mu^{\mathbf{Y}}_j)^2}{\lambda^{\mathbf{Y}}_j}\right| &\le \sum_{j=1}^{T_m}\frac{\left|(\bar{y}_j)^2-(\mu^{\mathbf{Y}}_j)^2\right|}{\lambda^{\mathbf{Y}}_j} + 2\sum_{j=1}^{T_m}\frac{\left|z_j(\bar{y}_j-\mu^{\mathbf{Y}}_j)\right|}{\lambda^{\mathbf{Y}}_j} \\
&= \sum_{j=1}^{T_m}\frac{\left|\bar{y}_j-\mu^{\mathbf{Y}}_j\right|\left|\bar{y}_j+\mu^{\mathbf{Y}}_j\right|}{\lambda^{\mathbf{Y}}_j} + 2\sum_{j=1}^{T_m}\frac{\left|z_j(\bar{y}_j-\mu^{\mathbf{Y}}_j)\right|}{\lambda^{\mathbf{Y}}_j}. \qquad (25)
\end{aligned}$$

If we take expectations and apply the Cauchy-Schwartz inequality, we have that

$$\mathbb{E}\left[\sum_{j=1}^{T_m}\frac{\left|\bar{y}_j-\mu_j^{\mathbf{Y}}\right|\left|\bar{y}_j+\mu_j^{\mathbf{Y}}\right|}{\lambda_j^{\mathbf{Y}}}\right] \leq \sum_{j=1}^{T_m}\left(\frac{\mathbb{E}\left[(\bar{y}_j-\mu_j^{\mathbf{Y}})^2\right]\mathbb{E}\left[(\bar{y}_j+\mu_j^{\mathbf{Y}})^2\right]}{(\lambda_j^{\mathbf{Y}})^2}\right)^{1/2}.$$

It happens that $\bar{y}_j-\mu_j^{\mathbf{Y}}$ and $\bar{y}_j+\mu_j^{\mathbf{Y}}$ are one-dimensional normal variables, with means equal to 0 and $2\mu_j^{\mathbf{Y}}$, respectively, and variances equal to $\lambda_j^{\mathbf{Y}}/m$ for $1\leq j\leq R_m^{\mathbf{Y}}$. Taking this into account, applying Jensen's inequality and the fact that $L_\mu^{\mathbf{X},\mathbf{Y}}<\infty$, we have

$$\begin{aligned}
\mathbb{E}\left[\sum_{j=1}^{T_m}\frac{\left|\bar{y}_j-\mu_j^{\mathbf{Y}}\right|\left|\bar{y}_j+\mu_j^{\mathbf{Y}}\right|}{\lambda_j^{\mathbf{Y}}}\right] &\leq \sum_{j=1}^{T_m}\left(\frac{1}{m}\left(\frac{1}{m}+\frac{(2\mu_j^{\mathbf{Y}})^2}{\lambda_j^{\mathbf{Y}}}\right)\right)^{1/2}\\
&\leq \left(\frac{T_m}{m}\sum_{j=1}^{T_m}\left(\frac{1}{m}+\frac{(2\mu_j^{\mathbf{Y}})^2}{\lambda_j^{\mathbf{Y}}}\right)\right)^{1/2}\\
&= \left(\frac{T_m}{m}\left(\frac{T_m}{m}+4\left\|(\Sigma_{T_m}^{\mathbf{Y}})^{-1/2}(\mu^{\mathbf{Y}}-\mu^{\mathbf{X}})_{T_m}\right\|^2\right)\right)^{1/2}\\
&= \left(\frac{T_m}{m}\right)^{1/2}\left(2\left\|(\Sigma_{T_m}^{\mathbf{Y}})^{-1/2}(\mu^{\mathbf{Y}}-\mu^{\mathbf{X}})_{T_m}\right\|+o(1)\right)\\
&= O(T_m m^{-1/2})=O(m^{-1/6}).
\end{aligned}$$

Now, let us consider the expectation of the second term in (25). Given $d\in\mathbb{N}$, let us denote $S_d=(\Sigma_d^{\mathbf{Y}})^{-1/2}\Sigma_d^{\mathbf{X}}(\Sigma_d^{\mathbf{Y}})^{-1/2}$. Having in mind that $\bar{y}_j$ is a normal random variable with mean $\mu_j^{\mathbf{Y}}$ and variance equal to $\lambda_j^{\mathbf{Y}}/m$ for $1\leq j\leq R_m^{\mathbf{Y}}$, applying the Cauchy-Schwartz and Jensen inequalities, and using the facts that $L_S^{\mathbf{X},\mathbf{Y}}<\infty$ and that $\mathbb{P}_{\mathbf{Z}}=\mathbb{P}_{\mathbf{X}}$, we obtain

$$\begin{aligned}
\mathbb{E}\left[\sum_{j=1}^{T_m}\frac{\left|z_j(\bar{y}_j-\mu_j^{\mathbf{Y}})\right|}{\lambda_j^{\mathbf{Y}}}\right] &\leq \sum_{j=1}^{T_m}\left(\left[\frac{\mathbb{E}[(z_j)^2]\mathbb{E}[(\bar{y}_j-\mu_j^{\mathbf{Y}})^2]}{(\lambda_j^{\mathbf{Y}})^2}\right]\right)^{1/2}\\
&= \frac{1}{m^{1/2}}\sum_{j=1}^{T_m}\left(\frac{\mathbb{E}[(z_j)^2]}{\lambda_j^{\mathbf{Y}}}\right)^{1/2}\\
&\leq \left(\frac{T_m}{m}\sum_{j=1}^{T_m}\frac{\mathbb{E}[(z_j)^2]}{\lambda_j^{\mathbf{Y}}}\right)^{1/2}\\
&= \left(\frac{T_m}{m}\mathrm{Trace}(S_{T_m})\right)^{1/2}=O(T_m m^{-1/2})=O(m^{-1/6}).
\end{aligned}$$

Therefore, both terms in equation (25) are $O_P(m^{-1/6})$. Since $R_m^{\mathbf{Y}}\xrightarrow{P}\infty$, we have that $\left\{\frac{1}{R_m^{\mathbf{Y}}}\sum_{i=1}^{T_m}\left|\frac{(z_j-\bar{y}_j)^2-(z_j-\mu_j^{\mathbf{Y}})^2}{\lambda_j^{\mathbf{Y}}}\right|\right\}_{m\geq 1}$ converges to zero in probability as $m\to\infty$.

To finalize, let us consider the case when $L_S^{\mathbf{X},\mathbf{Y}}$, or $L_\mu^{\mathbf{X},\mathbf{Y}}$ is infinite. Our problem is to show that if $\mathbb{P}_{\mathbf{Z}} = \mathbb{P}_{\mathbf{X}}$, then

$$\frac{1}{R_m^{\mathbf{Y}}} \sum_{j=1}^{R_m^{\mathbf{Y}}} \frac{(z_j - \overline{y}_j)^2}{\lambda_j^{\mathbf{Y}}} \xrightarrow{P} \infty \text{ as } m \to \infty. \tag{26}$$

This case is very similar to the last part of the proof of Proposition 7.1. Here, we have

$$\sum_{j=1}^{R_m^{\mathbf{Y}}} \frac{(z_j - \overline{y}_j)^2}{\lambda_j^{\mathbf{Y}}} = \|(\Sigma_{R_m^{\mathbf{Y}}}^{\mathbf{Y}})^{-1/2}(\mathbf{z} - \overline{\mathbf{y}}_m)\|^2.$$

Thus, if we denote $\mathbf{m}_d^{\mathbf{Y}} = (\Sigma_d^{\mathbf{Y}})^{-1/2}\mu^{\mathbf{Y}}$, $\mathbf{u}_d^{\mathbf{Y}} = (\Sigma_d^{\mathbf{Y}})^{-1/2}\mathbf{z}$ and $\bar{\mathbf{y}}_m^* = (\Sigma_{R_m^{\mathbf{Y}}}^{\mathbf{Y}})^{-1/2}(\bar{\mathbf{y}}_m - \mu^{\mathbf{Y}})$, then

$$\sum_{j=1}^{R_m^{\mathbf{Y}}} \frac{(z_j - \overline{y}_j)^2}{\lambda_j^{\mathbf{Y}}} = \|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2 + \|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2 + \|\bar{\mathbf{y}}_m^*\|^2 - 2\langle \mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}, \mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\rangle + 2\langle \bar{\mathbf{y}}_m^*, \mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}} - \mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\rangle. \tag{27}$$

By assumption (3), we know that

$$\lim_d \frac{1}{d}\|\mathbf{m}_d^{\mathbf{Y}}\|^2 = L_\mu^{\mathbf{X},\mathbf{Y}}.$$

If we assume that $L_\mu^{\mathbf{X},\mathbf{Y}} = \infty$, given $M > 0$, there exists $D > 0$ such that $d^{-1}\|\mathbf{m}_d\|^2 > M$ for every $d \geq D$. Thus, from (17), we have that

$$\mathbb{P}\left[\frac{1}{R_m^{\mathbf{Y}}}\|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2 \leq M\right] \leq \mathbb{P}[R_m^{\mathbf{Y}} < D] \to 0.$$

Concerning the second term in (27), it happens that the random vector $\mathbf{z}$ is independent from the sequence $\{R_m^{\mathbf{Y}}\}$. Thus, conditionally to this sequence, the distribution of the sequence $\{\|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2\}$ coincides with that of a subsequence of $\{\|\mathbf{u}_d^{\mathbf{Y}}\|^2\}$. However, along the proof of Proposition 7.1, we proved that

$$\frac{1}{d}\|\mathbf{u}_d^{\mathbf{Y}}\|^2 \xrightarrow{P} L_S^{\mathbf{X},\mathbf{Y}} \text{ as } d \to \infty.$$

From here, a proof similar to that one we developed for the sequence $\{\|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2\}$ allows us to conclude that if $L_S^{\mathbf{X},\mathbf{Y}} = \infty$, then

$$\frac{1}{R_m^{\mathbf{Y}}}\|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2 \xrightarrow{P} \infty, \text{ as } m \to \infty.$$

The same reasoning we employed in Proposition 7.1 is enough to prove that

$$\frac{1}{R_m^{\mathbf{Y}}} \frac{\langle \mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}, \mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\rangle}{\max(\|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2\|, \|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|^2)} \xrightarrow{P} 0 \text{ as } m \to \infty.$$

The distribution of $\|\bar{\mathbf{y}}_m^*\|^2$ is equal to that of a sum of $R_m^{\mathbf{Y}}$ squares of centered, one-dimensional normal variables with variance equal to m^{-1}. Thus, taking again $\{T_m\} \subset \mathbb{N}$ going to infinity at exact rate $m^{2/6}$, we would have that for every $\epsilon > 0$,

$$\mathbb{P}[\|\bar{\mathbf{y}}_m^*\|^2 > \epsilon] \leq \mathbb{P}[R_m^{\mathbf{Y}} > T_m] + \frac{1}{\epsilon} m^{-4/6}, \tag{28}$$

and consequently, $\|\bar{\mathbf{y}}_m^*\| \xrightarrow{P} 0$ as $m \to \infty$.

Only the last term in (27) remains to be analyzed. Here, the Cauchy-Schwartz inequality and equation (28) allow us to conclude that

$$\left|\langle \bar{\mathbf{y}}_m^*, \mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}} + \mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}} \rangle\right| \leq \|\bar{\mathbf{y}}_m^*\| \left(\|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\| + \|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\| \right) = o_P\left(\sup(\|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|, \|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|)^2 \right).$$

This proves that the leading terms in equation (27) are the first two. Thus, the fact that at least one of the sequences $\{\|\mathbf{m}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|\}$, or $\{\|\mathbf{u}_{R_m^{\mathbf{Y}}}^{\mathbf{Y}}\|\}$ goes to infinity gives (26). •

Proof of Proposition 4.2: Fix n (and, consequently, m_n is also fixed). Let us define

$$\eta(d_n) = \inf\{\hat{\lambda}_j^{\mathbf{X}} - \hat{\lambda}_{j+1}^{\mathbf{X}}, \hat{\lambda}_j^{\mathbf{Y}} - \hat{\lambda}_{j+1}^{\mathbf{Y}}, : j = 1, \ldots, d_n\}.$$

Obtain $S_n^{\mathbf{X}} \in \mathbb{N}$ (respectively, $S_{m_n}^{\mathbf{Y}} \in \mathbb{N}$) by applying (7) with $\eta(d_n)$ and the set of eigenvalues $\{\hat{\lambda}_1^{\mathbf{X}}, \hat{\lambda}_2^{\mathbf{X}}, \ldots\}$ (respectively, $\{\hat{\lambda}_1^{\mathbf{Y}}, \hat{\lambda}_2^{\mathbf{Y}}, \ldots\}$). By definition, $d_n \geq u_n$. So, we obviously have that $d_n \to \infty$, and hence $\eta(d_n) \to 0$. Moreover, the fact that $d_n \leq \min(R_n^{\mathbf{X}}, R_{m_n}^{\mathbf{Y}})$ implies that $\eta(d_n) \geq \eta_n^{\mathbf{X}}$ and $\eta(d_n) \geq \eta_{m_n}^{\mathbf{Y}}$. Consequently, we also have that $\lim_n n^{1/6}\eta(d_n) = \lim_n m_n^{1/6}\eta(d_n) = \infty$.

On the other hand, $d_n \leq S_n^{\mathbf{X}}$, and then it is easy to show that d_n satisfies (17) with $\eta_n^{\mathbf{X}}$ replaced by $\eta(d_n)$. So, if $j \leq d_n$, then statements (14), (15), (18) and (19) also hold (with $\eta_n^{\mathbf{X}}$ replaced by $\eta(d_n)$ whenever required). From this point, it is possible to repeat the proof of Theorem 4.1 to obtain that if $\mathbf{z}$ is an observation produced by $\mathbb{P}_{\mathbf{X}}$, then

$$|\hat{D}_{d_n}^{\mathbf{Y}}(\mathbf{z}, \bar{\mathbf{y}}_{m_n}) - 1| - |\hat{D}_{d_n}^{\mathbf{X}}(\mathbf{z}, \bar{\mathbf{x}}_n) - 1| \xrightarrow{P} |L_\mu^{\mathbf{X},\mathbf{Y}} + L_S^{\mathbf{X},\mathbf{Y}} - 1| > 0$$

as $n \to \infty$, and therefore, $\mathbf{z}$ is asymptotically correctly classified. The same arguments with $S_{m_n}^{\mathbf{Y}}$ gives the result when the observations are produced from $\mathbb{P}_{\mathbf{Y}}$. •

References

[1] Baíllo, A., Cuevas, A. and Cuesta-Albertos, J. A. (2011) Supervised classification for a family of Gaussian functional models. *Scandinavian J. Statist.*, Vol. 38, 480-498.

[2] Berrendero, J. R., Cuevas, A. and Torrecilla, J. L. (2015) On near perfect classification and functional Fisher rules via reproducing kernels. *Available at arXiv:1507.04398v1.*

[3] Delaigle, A. and Hall, P. (2012) Achieving near perfect classification for functional data. *J. R. Statist. Soc. B*, Vol. 74, 267-286.

[4] Ferraty, F. and Vieu, P. (2006) *Nonparametric Functional Data Analysis.* New York: Springer.

[5] Feldman, J. (1958) Equivalence and perpendicularity of Gaussian processes. *Pacific J. Math.*, Vol. 8, No. 4, 699-708.

[6] Galeano, P., Joseph, E. and Lillo, R. E. (2015) The Mahalanobis distance for functional data with applications to classification. *Technometrics*, Vol. 57, 281-291.

[7] Genz, A., Bretz, F., Miwa, T., Mi, X., Leisch, F., Scheipl, F. and Hothorn, T. (2015) mvtnorm: Multivariate normal and t distributions. R package version 1.0-3.

[8] Hajek, J. (1958). A property of J-divergences of marginal probability distributions. *Cz. Math. J.*, Vol. 8, 460-462.

[9] Hall, P. and Hosseini-Nasab, M. (2006) On properties of functional principal components analysis. *J. R. Statist. Soc. B*, Vol. 68, 109-126.

[10] Hastie, T., Tibshirani, R. and Friedman, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer.

[11] Iacus, S. M. (2015) sde: Simulation and inference for stochastic differential equations. R package version 2.0.14.

[12] Rao, C. R. and Varadarajan, V. S. (1963) Discrimination of Gaussian processes. *Sankhya, Series A*, Vol. 25, 303-330.

[13] Sato, H. (1967) On the equivalence of Gaussian measures. *J. Math. Soc. Japan*, Vol. 37, 321-354.

[14] Shepp, L. A. (1966) Radon-Nikodym derivatives of Gaussian measures. *Ann. Math. Statist.*, Vol. 37, 321-354.

[15] Shepp, L. A. (1966) Gaussian measures in function space. *Pacific J. Math.*, Vol. 17, 167-173.

[16] Varberg, D.E. (1961) On equivalence of Gaussian measures. *Pacific J. Math.*, Vol. 11, 751-762.

[17] Zhang, J.-T., Liang, X. and Xiao, S. (2011) On the two-sample Behrens-Fisher problem for functional data. *J. Statist. Theo. Pract.*, Vol. 4, 571-587.